

HTAR Reference Manual

Table of Contents

Preface	3
Introduction	4
HTAR's Role	4
How HTAR Works	6
TAR and HTAR Compared	8
How to Use HTAR	10
HTAR Execute Line	10
HTAR Error Conditions	13
HTAR Limitations and Restrictions	15
HTAR Environment Variables	17
HTAR Options	18
Action Options	18
Archive Option	19
Control Options	20
HTAR Examples	25
Creating an HTAR Archive File	25
Retrieving Files from within an Archive	27
Rebuilding a Missing Index	29
Specifying Very Many Files	31
Archiving Between Nonstorage Machines	32
Disclaimer	34
Keyword Index	35
Alphabetical List of Keywords	36
Date and Revisions	37

Preface

- Scope:** This HTAR Reference Manual explains the roles, usage, options, and features of LC's locally developed file-bundling and storage utility (the "HPSS Tape Archiver" or HTAR). Besides introducing the HTAR execute line, control options, and environment variables, the text also compares HTAR with standard TAR and provides annotated examples of using HTAR for the special tasks it is designed to handle (such as retrieving files from within a stored archive, successfully managing very large archives, or depositing an archive in a designated nonstorage location).
- Availability:** HTAR runs on all LC production IBM (AIX), Compaq (Tru64), and Linux/CHAOS machines, on both open and secure graphics machines, and on many LC special-purpose Suns.
- Consultant:** For help contact the LC customer service and support hotline at 925-422-4531 (open e-mail: lc-hotline@llnl.gov, SCF e-mail: lc-hotline@pop.llnl.gov).
- Printing:** The print file for this document can be found at:
- on the OCF: <http://www.llnl.gov/LCdocs/htar/htar.pdf>
on the SCF: https://lc.llnl.gov/LCdocs/htar/htar_scf.pdf

Introduction

HTAR's Role

GOALS.

HTAR ("HPSS Tape Archiver") is an LC-designed TAR-like utility program that makes TAR-compatible archive (library) files but with storage support and enhanced archive-management features. Despite its misleading name, HTAR does *not* write files to tape, but to LC's open or secure archival storage (HPSS) disk farm (or, optionally, to other specified LC hosts).

HTAR's enhancements include its ability to:

- bundle many small files together in memory (without using more local disk space) for more efficient handling and transfer,
- send the resulting large archive file directly to (open or secure) storage without your needing to invoke FTP separately (or to another LC machine if you use -F),
- retrieve individual files from a stored archive *without* moving the whole large archive back to your local machine first (or, optionally, without even staging the whole archive to disk), and
- accelerate transfers to and from storage by deploying multiple threads and by using as many parallel interfaces to storage as are available on the (production) machine where it runs.

SCOPE.

A subsection of this introduction (below (page 8)) compares traditional UNIX TAR with LC's enhanced HTAR feature by feature to reveal the value of this added tool. But in general HTAR maintains full output compatibility with the POSIX 1003.1 standard TAR format while successfully archiving hundreds or even thousands of incoming files, handling files of greatly mixed sizes or types, and imposing no limit on the total size of the archive files that it builds.

PERFORMANCE.

HTAR is designed to run quickly, especially when transferring large archives to storage. Its use of multiple concurrent threads and HPSS parallel disk striping enable HTAR to reach transfer rates to storage as high as 150 Mbyte/s, which exceeds 30 times the normal rate for transferring many small (around 4-Mbyte) files separately. Users concerned about to-storage transfer rates should use the open or secure NETMON (URL: <http://www.llnl.gov/LCdocs/netmon>) web sites to list or plot the latest to-storage network statistics.

FILE MANAGEMENT.

HTAR can store archive-member files as large as 8 Gbyte. There is no maximum size for a whole HTAR archive (some have successfully reached 200 Gbyte). HTAR makes single copies of each stored archive (COS 140) by default, but you can request dual-copy storage (for extra safety) of a mission critical archive of any size by using HTAR's -Y 150 option.

THIS GUIDE.

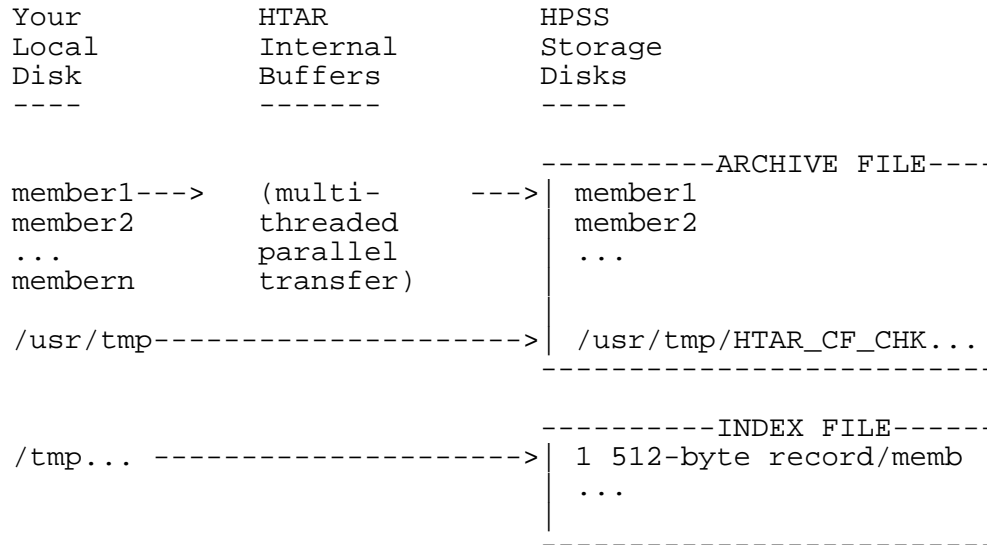
Because HTAR combines two features usually separate (file bundling and file storage), having some understanding of how it works can help you use it effectively. So one subsection below shows and explains the relationship among the three files (the archive file, the index file, and the consistency file) that HTAR uses to manage every transaction. Also shown is a feature-by-feature comparison of HTAR with traditional UNIX TAR.

One major section in this manual tells how to run HTAR, and spells out common error conditions, known limitations (with work-arounds), and HTAR environment variables. A second major section describes the function of each HTAR option (distinguishing the required action options from the control options). A third section gives annotated step-by-step examples of how to use HTAR to handle common file-archiving tasks and problems (including creation of archives containing *very* many files and optional transfer of archives to or from *nonstorage* LC machines).

HTAR users may also benefit from familiarity with another LC-developed specialty tool that provides nonstandard file-handling and file-transfer features linked to file storage, namely NFT (see the [NFT Reference Manual](http://www.llnl.gov/LCdocs/nft) (URL: <http://www.llnl.gov/LCdocs/nft>) for details). HTAR itself does not, however, use NFT's "persistence" mechanisms to manage file-storage delays. For a general introduction to LC storage tools and techniques, see [EZSTORAGE](http://www.llnl.gov/LCdocs/ezstorage). (URL: <http://www.llnl.gov/LCdocs/ezstorage>)

How HTAR Works

HTAR makes an archive (or library) file in the standard POSIX 1003.1 TAR format, which allows TAR to open any HTAR archive file. But HTAR offers more services than ordinary TAR, and it therefore needs extra internal machinery to support those services. While much of this extra machinery is hidden during normal use, some of it reveals itself in HTAR status messages or command responses that might prove surprising or confusing without some insight into how HTAR works. So this section briefly explains how HTAR makes an archive file and the role that several support files play in that process.



Archive File (*name.tar*)

When you run HTAR with the create-archive (-c) option, the program first opens a connection to storage (HPSS). It then deploys multiple threads to transfer in parallel the local-disk files that you specify (the "archive members," left side of the diagram) into a TAR-format envelope file created (unless you request otherwise) in your storage home directory (right side in the diagram). This archive file never exists on local disk (unless you demand it with the -E option), even in temporary directories on the machine where HTAR runs. Instead, HTAR reads the member files piecewise into its internal buffers and moves the data directly to HPSS (or to another host specified by -F), where it assembles the archive.

HTAR simultaneously builds a separate index file (outside the archive) and a little consistency file (deposited last inside the archive), discussed below. HPSS is very reliable, but HTAR automatically uses a storage "class of service" (COS 140) that keeps only one copy of your stored archive file. For files of special importance (only), use HTAR's -Y 150 option to force creation of a duplicate (invisible) backup copy.

Index File (*name.tar.idx*)

To allow archives of unlimited size and to support the direct extraction of any stored archive member(s) *without* retrieving the whole archive to local disk, HTAR automatically builds an external index file to accompany every archive that you create. While making the archive, HTAR temporarily writes the index file to the local /tmp file system on the machine where it runs (left side of the diagram), then transfers it (by default) to the same storage (or other remote) directory where the archive itself resides at the end of the process (right side of the diagram).

Each HTAR index file contains one 512-byte record for every member file stored in the corresponding archive file, regardless of the member file's size (so even a 10,000-file archive will have an index file of only about 5 Mbyte). HTAR index files are so much smaller than the archives that they support that the index file often remains on HPSS disk (to rapidly respond to queries) even when the larger archive file itself migrates to storage tape (cartridges). If you use HTAR's -E or -F options to force the archive to a location other than storage, the index file goes to the same location as the archive file.

Consistency File (/usr/tmp/HTAR_CF_CHK_#####)

Because the archive and index files are separate, HTAR maintains a consistency check between them in an additional 1-block (256-byte) file always parked (as a last step) at the end of each archive. This consistency file's name has the long numerical format shown above, but begins with /var/tmp if generated on a Compaq machine and /tmp if generated on a Sun. HTAR never extracts this file (unless you specifically request it), but every use of -t and -v (together with -c or -x) reports this perhaps unexpected consistency file at the end of HTAR's list of archived contents.

TAR and HTAR Compared

TAR was originally intended to write a specified set of files to offline tape (or retrieve them), and hence by extension, to simply write (or retrieve) a specified set of files to a local envelope or library file for easier management. HTAR in many ways returns TAR to its roots because it is specifically designed to efficiently *store* a set of files together in HPSS or get them back, not merely to make an archive file and leave it, although you can force HTAR to do that.

This table compares the more familiar TAR features and effects with those (often enhanced) of HTAR:

Feature	TAR	HTAR
Can create an archive file without storing it?	Yes (the default)	Only with -E or -F
Can create an archive file without using local disk space?	No	Yes (the default)
Can store an archive file while creating it?	No, needs FTP	Yes (the default)
Can store an archive file without creating it?	No	No
Can write an archive to (offline) tape?	Yes	No (to storage disk)
Can write an archive to another machine?	No	Yes, with -F
Can read an archive from another machine?	No	Yes, with -F
Can read any TAR archive file?	Yes (the default)	Yes, if -X first
Can read any HTAR archive file?	Yes	Yes (the default)
Can extract just one file from a <i>stored</i> archive?	No	Yes (the default)
Can add file(s) to an existing archive?	Yes	No
Default target if no archive specified?	Yes (tape)	No, -f required
Treats input directories recursively?	Yes	Yes (-L disables)
Preserves original permissions on files?	No (uses UMASK)	Yes, with -p
Depends on HPSS availability to work?	No	Yes

Feature	TAR	HTAR
Archive duplicated automatically in storage?	No	Only with -Y 150
Builds and needs an external index file?	No	Yes
Builds and needs a consistency check file?	No	Yes
Overwrites existing files without warning?	Yes	Yes (-w disables)
Can use standard input or output?	Yes (with -f -)	No
Order of options important?	No	Yes
Table of contents (-t) reveals what?	File names only	File names and properties

How to Use HTAR

HTAR Execute Line

To run HTAR you must log on to an LC production machine where HTAR has been installed at a time when the storage system (HPSS) is up and available to users. The HTAR execute line has the general form

```
htar action archive [options] [filelist]
```

and the specific form

```
htar -c|t|x|X -f archivename [-BdEFhHILmMopSTvVwY] [flist]
```

where exactly one action and the *archivename* are always required, while the control options and (except when using -c) the *filelist* (or *flist*) can be omitted (and the options can share a hyphen flag with the action for convenience). Here,

- c (create) opens a parallel FTP connection to storage, creates an archive file at the storage location (*not* online) and with the name specified by -f, and transfers (copies) into the archive each file specified by *filelist* (required whenever you use -c). If *archivename* already exists, HTAR overwrites it without warning. To create a local archive file instead (the way TAR does), also use -E; to deposit it on a *nonstorage* host, also use -F. If *filelist* specifies any directories, HTAR includes them and all of their children recursively (using -L disables this recursion).
- t (table of contents) opens a parallel FTP connection to storage, then lists (in the order that they were stored) the files currently within the stored archive file specified by -f, along with their owner, size, permissions, and modification date (the list includes HTAR's own consistency file (page 6)). Here *filelist* defaults to * (all files in the archive), but you can specify a more restrictive subset (usually by making *filelist* a filter).
- x (lowercase eks, extract) opens a parallel FTP connection to storage (or to another remote host specified by -F), then transfers (extracts, copies) from the stored (remote) archive file specified by -f each internal file specified by *filelist* (or all files in the archive if you omit *filelist*). If *filelist* specifies any directories, HTAR extracts them and all their children recursively (using -L disables recursion). If any file already exists locally, HTAR overwrites it without warning, and it creates all new files with the same owner and group IDs (and if you use -p, with the same UNIX permissions) as they had when stored in the archive. (If you lack needed permissions, extracted files get your own user and group IDs and the local UMASK permissions; if you lack write permission then -x creates no files at all.) Note that -x works directly on the remote archive file; you never retrieve the whole archive from storage just to extract a few specified files from within it (impossible with TAR).

-X (uppercase eks, index) opens a parallel FTP connection to storage (or to another remote host specified by **-F**), then creates an (external) index file for the existing archive file specified by **-f** (a stored TAR-format file by default, a local TAR-format file if you also use **-E**). Using **-X** rescues an HTAR archive whose (stored) index file was lost, and it enables HTAR to manage an archive originally created by traditional TAR. The resulting external index file is stored if the corresponding archive is stored, but local if the archive is local (with **-E**). See the "[How HTAR Works](#) (page 6)" section for an explanation of HTAR index files.

-f *archivename*

(required option) specifies the archive file on which HTAR performs the **-c|t|x|X** actions described above. HTAR has no default for **-f** (whose argument must appear immediately after the option name). Since HTAR (normally) operates on *stored* archive files, *archivename* also locates the archive file relative to your *storage* (HPSS, not online) home directory: a simple file name here (e.g., *abc.tar*) resides in your storage home directory, while a relative pathname (e.g., *xyz/abc.tar*) specifies a subdirectory of your storage home directory (i.e., */users/unn/username/xyz/abc.tar*), the recommended practice for batch jobs. Never use tilde (*~*) in *archivename* because the shell expands it into your online, not your storage, home directory. HTAR's **-f** makes no subdirectories; you must have created them in advance (with FTP's *mkdir* option) before you mention them in *archivename*. When used with **-F** to make or read an archive on a nonstorage machine, *archivename* should be the full pathname of the archive on the remote machine (e.g., */var/tmp/abc.tar*).

filelist specifies the input files for **-c** and the subset of archived files to process (for **-t**, **-x**, or **-X**). Omitting *filelist* for **-c** yields a null result and the error message "refusing to create empty archive." Omitting *filelist* for **-t**, **-x**, or **-X** defaults to ***, all files within the archive file specified by **-f**. Here *filelist* can include a blank-delimited list of files, UNIX file filters (metacharacters), or directory name(s) to be processed recursively. Use **-L** to specify directories without recursion.

SYNTAX ISSUES.

Traditional TAR allows great flexibility in ordering its options or combining them with one option flag, so that

```
tar -vfc abc.tar *
```

works just the same as

```
tar -c -v -f abc.tar *
```

But HTAR requires more care in ordering options on its execute line, so that

```
htar -vfc abc.tar *
```

fails completely, yielding only a syntax warning reminder but no output. With HTAR:

(1) "actions" **-c|t|x|X** *must precede* all other options (required whether they share a single option flag (**-**) or each have their own), and

(2) the archive specifier **-f** must *immediately* precede its argument (the *archivename*), again regardless of whether **-f** has its own flag or shares with others.

Thus only these combinations

```
htar -cvf abc.tar *
```

```
htar -c -v -f abc.tar *
```

will work for the example case begun above.

DEFAULTS.

By default, HTAR creates an archive by copying files from the online directory where you run it into a file in your storage (HPSS) home directory, and it extracts files by reversing that process. You must always specify the name of the archive file on which HTAR operates (there is never a default archive). Once you name the archive, HTAR calls the corresponding external index file *archivename.idx* by default and stores it in the same HPSS directory as the archive (by default). HTAR's -I option lets you specify a nondefault name or location for the index file. The HTAR consistency file's name begins with /usr/tmp/HTAR on IBM machines, /var/tmp/HTAR on Compaq (Tru64) machines, and /tmp/HTAR on Sun machines (a default that you cannot change). On Linux/CHAOS machines, the consistency file's name begins with /var/tmp/*uname*/HTAR, where *uname* is your login name on the machine where you run HTAR. By default HTAR stores a single copy of each archive in HPSS (COS 140), regardless of its size. But you can request dual-copy storage (COS 150) of any mission critical HTAR archive, regardless of its size, by using the -Y 150 option on HTAR's execute line.

HTAR Error Conditions

The most common error conditions and HTAR's (often cryptic) responses to them are summarized here to help you troubleshoot:

Storage (HPSS) is down.

When HPSS is unavailable to users (perhaps for maintenance), no stored archive can be read or written. HTAR returns a message of this form and ends (there is no persistence as with NFT):

```
hpssex_OpenConnection: unable to obtain remote site info
result = -5000, errno = 0
Unable to setup communication to HPSS. Exiting...
###WARNING  htar returned non-zero exit status.
              255 = /usr/local/bin/htar.exe...
```

Specified archive directory does not exist.

If -f specifies a child directory (of your storage home directory) that you have *not* previously created (with FTP's mkdir option), HTAR returns an error message that often only hints at the actual problem. When you attempt to create an archive in a nonexistent (sub)directory, HTAR responds:

```
***Error -2 on hpss_Open (create) for archivename
###WARNING  htar returned non-zero exit status.
              72 = /usr/local/bin/htar.exe...
```

When you attempt to extract files from an archive in a nonexistent (sub)directory, HTAR at least replaces the first line of this error message with:

```
***Fatal error opening index file archivename.idx
###WARNING  htar returned non-zero exit status.
              72 = /usr/local/bin/htar.exe...
```

Specified archive file does not exist.

If -f specifies an archive file that does not exist (perhaps because you deleted it or mistyped its name), HTAR responds:

```
[FATAL] no such HPSS archive file: archivename
###WARNING  htar returned non-zero exit status.
              72 = /usr/local/bin/htar.exe...
```

Specified index file does not exist.

If you try to list (-t) or extract (-x) files from an actual HTAR archive whose corresponding external index file (*archivename.idx*) has been deleted or moved, HTAR pinpoints the problem only by reporting the missing index name:

```
No such file: archivename.idx
###WARNING  htar returned non-zero exit status.
              72 = /usr/local/bin/htar.exe...
```

You can work around the missing index by using HTAR's -X (uppercase eks) option to rebuild the index while the archive remains stored, or you can retrieve the whole archive from storage with FTP and then open it with TAR.

HTAR's *filelist* omitted.

If you try to create (-c) an archive without specifying a *filelist* (or without using a *filelist* replacement such as -L), HTAR connects to HPSS but quickly ends with the message

```
Refusing to create empty archive.
```

If you try to list (-t) or extract (-x) without specifying a *filelist*, HTAR defaults to processing all files in the archive.

HTAR run with no options.

Because HTAR requires exactly one action (-c|t|x|X) and a specified archive file (-f) to run, executing the program with nothing else on the execute line yields a terse syntax summary. There is no prompt for input, and HTAR terminates.

Command line too long for shell.

The easy way to build an HTAR archive of *very many* like-named files is to specify them indirectly by using a UNIX metacharacter (filter, wild card) such as * (to match any string) or ? (to match any single character). But if the selected file set has thousands of members, the list of input names that the UNIX shell generates by expanding such an "ambiguous file reference" may grow too long to handle. See the Limitations and Restrictions (page 15) section below for several ways to work around such excessively long command lines when building large archives with HTAR.

Wild cards (metacharacters) used for retrieval.

HTAR allows * only to create an archive, not to retrieve files from one ("no match" is the usual, but not the only possible, error message). See the Retrieving Files (page 27) section below for an analysis and possible ways to work around this limitation.

HTAR Limitations and Restrictions

The current version of HTAR has the following known limitations or usage restrictions:

- **I/O:**
HTAR cannot read from UNIX pipes. For example, you cannot pipe in (with |) a list of file names generated on the fly by CAT or FIND (you must use HTAR's -L option instead). However, you *can* redirect HTAR output into a file for separate postprocessing (see the [Retrieving Files](#) (page 27) section for one application of this).
- **METACHARACTERS:**
HTAR leaves all processing of metacharacters (filters or wild cards, such as *) to the shell. This means that when you create an HTAR archive you can use * to select from among your local files to store, but when you *retrieve* specific files from within an already stored archive you CANNOT use * to select from among the stored files to get back. See the [Retrieving Files](#) (page 27) example below for details on this limitation, and a few suggested but inelegant ways to work around it. Another side effect of this approach to metacharacters is that C shell (CSH) users must type the three-character string -\? (instead of -?) to display HTAR's help message.
- **UPDATES:**
No options exist to update (replace), remove (delete), or append to individual files that HTAR has already archived. You must replace (create again) an entire archive to alter the member files within it.
- **NAME LENGTH:**
To comply with POSIX 1003.1 standards regarding TAR-file input names, the longest input file name of the form prefix/name that HTAR can accept has 154 characters in the prefix and 99 characters in the name. Link names likewise cannot exceed 99 characters.
- **FILE SIZE:**
The maximum size of a single member file within an HTAR archive is 8 Gbyte (a limit imposed by the format of the TAR header). HTAR imposes no limit on the maximum size of an archive file (some have successfully reached 200 Gbyte), but local disk space (when using -E or -F) or storage space might externally limit an archive's size. Users can specify a maximum number of member files per archive with HTAR's -M option.
- **PASSWORDLESS FTP:**
Because HTAR (unlike FTP) does not support user dialog with a server and has no password-passing option, you can only manipulate HTAR archives on machines with preauthenticated (passwordless) FTP servers. This limits the use of HTAR's -F option to LC production machines only, since there is no way to satisfy the password request from other FTP servers.

TOO MANY NAMES.

For users who make HTAR archives containing thousands of files, a different kind of limitation poses problems, a limitation of the UNIX shell (csh, bsh, ksh) rather than of the HTAR program itself. One would normally select multiple files for archiving by using a UNIX "ambiguous file reference," a partial file name adjacent to one or more shell metacharacters (or "wild card" filters, such as the asterisk(*)). Your current

shell automatically expands the metacharacter(s) to generate a (long) alphabetical list of matching file names, which it inserts into the execute line as if you had typed them all yourself. Thus

```
htar -cf test.tar a*
```

might become equivalent to a command line with dozens of a-named files on the end. Each shell has a maximum length for execute lines, however, and if your specified metacharacter filter matches thousands of file names, HTAR's execute line may grow too long for the shell to accept. This would prevent building your intended many-file archive.

WORK-AROUND 1: USE A DIRECTORY.

The most effective, least resource-intensive way to work around the problem of having a (virtual) HTAR execute line too long for the shell to handle is to plan ahead and keep (or generate) in a single directory all and only the files that you want to archive. HTAR processes directory names recursively by default. So if you specify only the relevant directory name on HTAR's execute line, HTAR will (internally) archive every file within the directory without any filter-induced length problems. For example,

```
htar -cf test.tar projdir
```

will successfully archive any number of files within the PROJDIR directory yet use no troublesome shell-mediated file-name generation to do it.

WORK-AROUND 2: USE FIND.

The UNIX FIND utility is designed to produce lists of files (that meet specified criteria) to feed into other programs for further processing. So FIND offers a second way for HTAR to archive very large numbers of files without having a very long execute line. Indirection is required for success, however. The "natural" use of FIND's -EXEC option to run another program (here, HTAR) driven by a list of files from FIND, for example,

[WRONG]

```
find . -name 'a*' -print -exec htar -cf test.tar {} \;
```

fails to produce the desired effect. This actually runs HTAR once (to build an archive called test.tar) for each successive input file (here, files beginning with A). If there are thousands of files, HTAR just repeatedly creates a one-file archive thousands of times (each replacing the previous archive) so that only the last file processed really remains in test.tar at the end.

Instead, you must build an intermediate, external list of file names (in a file) with FIND, then have HTAR process that list with its -L option. The correct sequence is:

```
find . -name 'a*' -print > alist
htar -cf test.tar -L alist
rm alist
```

(The use of the metacharacter * in FIND's execute line here does not pose the same too-long problem as it did originally in HTAR's execute line because the surrounding quotes shelter the filter from shell processing. FIND's -NAME option generates the list of matching names internally, without expanding FIND's execute line.) Later, if HTAR is modified to accept piped input, a UNIX pipe (|) between FIND and HTAR could eliminate the need for an overt external list-of-names file while using this same strategy.

HTAR Environment Variables

HTAR uses the following HPSS-related environment variables if they are available on the machine where it runs:

HPSS_HOSTNAME

specifies the host name or IP address of the network interface to which HPSS mover(s) should connect when transferring data at HTAR's request (overridden by the file specified in PFTP_CONFIG_FILENAME). The default interface (the alternative to HPSS_HOSTNAME) is often slow, such as the control Ethernet of an IBM SP machine.

HPSS_PATH_ETC

specifies the pathname of a local directory containing the HPSS network options file.

HPSS_SERVER_HOST

specifies the server host name and optional port number of the HTAR server.

HTAR_COS

specifies the default class of service (COS) ID for the archive file that HTAR creates, or contains the string AUTO to force HPSS to automatically select the class of service based on the file size. HTAR option -Y overrides HTAR_COS (see the end of the next section for details).

PFTP_CONFIG_FILENAME

specifies the pathname of a file (usually /etc/pftp_config) that contains the list of HPSS network interfaces to be used for parallel transfers of files to and from stored HTAR archives.

HTAR Options

Action Options

Exactly ONE of these action options is required every time that you run HTAR.

- c (create) opens a parallel FTP connection to storage, creates an archive file at the storage location (*not* online) and with the name specified by -f, and transfers (copies) into the archive each file specified by *filelist* (required whenever you use -c). If *archivename* already exists, HTAR overwrites it without warning. To create a local archive file instead (the way TAR does), also use -E; to deposit it on a *nonstorage* host, also use -F. If *filelist* specifies any directories, HTAR includes them and all of their children recursively (using -L disables this recursion).
- t (table of contents) opens a parallel FTP connection to storage, then lists (in the order that they were stored) the files currently within the stored archive file specified by -f, along with their owner, size, permissions, and modification date (the list includes HTAR's own consistency file (page 6)). Here *filelist* defaults to * (all files in the archive), but you can specify a more restrictive subset (usually by making *filelist* a filter).
- x (lowercase eks, extract) opens a parallel FTP connection to storage (or to another remote host specified by -F), then transfers (extracts, copies) from the stored (remote) archive file specified by -f each internal file specified by *filelist* (or all files in the archive if you omit *filelist*). If *filelist* specifies any directories, HTAR extracts them and all their children recursively (using -L disables recursion). If any file already exists locally, HTAR overwrites it without warning, and it creates all new files with the same owner and group IDs (and if you use -p, with the same UNIX permissions) as they had when stored in the archive. (If you lack needed permissions, extracted files get your own user and group IDs and the local UMASK permissions; if you lack write permission then -x creates no files at all.) Note that -x works directly on the remote archive file; you never retrieve the whole archive from storage just to extract a few specified files from within it (impossible with TAR).
- X (uppercase eks, index) opens a parallel FTP connection to storage (or to another remote host specified by -F), then creates an (external) index file for the existing archive file specified by -f (a stored TAR-format file by default, a local TAR-format file if you also use -E). Using -X rescues an HTAR archive whose (stored) index file was lost, and it enables HTAR to manage an archive originally created by traditional TAR. The resulting external index file is stored if the corresponding archive is stored, but local if the archive is local (with -E). See the "How HTAR Works (page 6)" section for an explanation of HTAR index files.

Archive Option

This option is required every time that you run HTAR.

-f *archivename*

(required option) specifies the archive file on which HTAR performs the -c|t|x|X actions described above. HTAR has no default for -f (whose argument must appear immediately after the option name). Since HTAR (normally) operates on *stored* archive files, *archivename* also locates the archive file relative to your *storage* (HPSS, not online) home directory: a simple file name here (e.g., abc.tar) resides in your storage home directory, while a relative pathname (e.g., xyz/abc.tar) specifies a subdirectory of your storage home directory (i.e., /users/unn/username/xyz/abc.tar), the recommended practice for batch jobs. Never use tilde (~) in *archivename* because the shell expands it into your online, not your storage, home directory. HTAR's -f makes no subdirectories; you must have created them in advance (with FTP's mkdir option) before you mention them in *archivename*. When used with -F to make or read an archive on a nonstorage machine, *archivename* should be the full pathname of the archive on the remote machine (e.g., /var/tmp/abc.tar).

Control Options

These options change how HTAR behaves, but none is required (default values are indicated when they exist).

- ? displays a short help message (a syntax summary of the HTAR execute line and a one-line description of each option). Users running HTAR under the C shell (CSH) will probably have to use the three-character string `-\'?` to display this help message.
- B adds block numbers to the listing (`-t`) output (normally used only for debugging).
- d *debuglevel*
(default is 0) sets to an integer from 0 through 5 the level of debug output from HTAR, where 0 disables debug information for normal use and 1 to 5 enable progressively more elaborate debug output.
- E emulates TAR by forcing the archive file to reside on the *local* machine (where you run HTAR) rather than in HPSS (storage), where it resides by default (`-f` always specifies the archive pathname, which `-E` interprets as local rather than remote). See also `-F` for making nonstorage *remote* archives. The HTAR index file goes into the same (local) directory as the archive.
- F [*user@*]*host*[*#port*]
overrides the HTAR default of a *stored* archive and specifies on which remote machine (*host*) the archive resides other than in HPSS. For creating archives, *host* is the sink machine; for extracting files from existing archives *host* is the source machine (see the between-machine [example](#) (page 32) for how to use `-F` properly). See also `-E` for making nonstorage *local* archives. The HTAR index file goes into the same (remote) directory as the archive. Any LC production machine with preauthenticated (passwordless) FTP service can be the `-F host`. HTAR still contacts the HPSS server even though the archive does not reside in HPSS, just to log all `-F` transactions. The *user* and *port* fields are seldom needed or appropriate because they usually betray the need for an FTP password and HTAR has no means to transmit one (the default *user* is you, the default *port* is 21).
- h (used only with `-c`; has no effect otherwise) for each symbolic link that it encounters, causes HTAR to replace the link with the actual contents of the linked-to file (stored under the link name, not under the file's original name). Later use of `-t` or `-x` treats the linked-to file as if it had always been present as an actual file with the link name. Without `-h`, HTAR records, reports, and restores every symbolic link overtly, but it does *not* replace the link with the linked-to contents.
- H *subopt[:subopt...]*
specifies a colon-delimited list of HTAR suboptions to control program execution. Possible *subopt* values now include:

cksum=on off	(default is off) enables or disables generating checksums when copying member files into the archive. Enabling checksums usually degrades HTAR's I/O performance and increases its CPU utilization.
nostage	avoids prestaging tape-resident (stored) archive files when HTAR performs -x, -X, or -t actions.
rmlocal	removes local member files after HTAR successfully writes them into the archive file (used with -c).
verify= test[test ,]	performs posttransfer verification after creating an archive, where <i>test</i> can be any of:
cksum ncksum	enables or disables (the default) verifying member file checksums by reading the archive file or by comparing the index file checksum with a checksum of the local member files.
compare nocompare	enables or disables (the default) byte-by-byte comparison of the local member files with the corresponding archive files.
paranoid noparanoid	enables or disables (the default) extreme efforts to detect problems (such as discovering whether local files were modified during archive creation before deleting them if authorized by RMLOCAL).

-I *indexname*

specifies a nondefault name for the HTAR external index file that supports the archive specified by -f.

WARNING: if you use -I to make any nondefault index name (3 cases, below) when you create (-c) an archive, then you **MUST** also use -I with the same argument every time you extract (-x) files from that archive (else HTAR will look for the default index, not find it, and end with an error).

There are three cases based on the first character of *indexname*:

- . (dot) If *indexname* begins with a period (dot), HTAR treats it as a suffix to append to the current archive name.
Example: -I .xnd yields an index file called *archivename.xnd*
- / If *indexname* begins with a / (slash), HTAR treats it as an absolute pathname (you must create all the subdirectories ahead of time with FTP's mkdir option).
Example: -I /users/unn/*yourname*/projects/text.idx uses that absolute pathname in storage (HPSS) for the index file.
- other* If *indexname* begins with any other character, HTAR treats it as a relative pathname (relative to the storage directory where the archive file resides, which might be different than your storage home directory).
Example: -I projects/first.index locates first.index at *storagehome*/projects/first.index if the archive file is in your *storagehome* (the default), but tries to locate first.index at *storagehome*/projects/projects/first.index if the archive was specified as -f projects/*aname* in the first place. (All such subdirectories must be created in advance.)
- L *inputfile* (used with -c) writes the files and directories specified by their literal names (in the *inputfile*, which contains file names one per line) into the archive specified by -f. Directories are *not* treated recursively; only a directory entry (but not its subdirectories or subfiles) are written to the archive. Normal metacharacters (tilde, asterisk, question mark) are treated literally, not expanded as filters.
(used with -x) retrieves the files and directories specified by their literal names. See the [Retrieving Files](#) (page 27) example below for how to use -L instead of wild cards to retrieve only specified files from a stored archive.
- m (used only with -x) makes the time of extraction the last-modified time for each member file (the default preserves each file's original time of last modification). The timestamp on directories is unpredictable because (1) the operating system may update the directory's last-modified time whenever HTAR extracts files from it, and (2) HTAR with -m will update the timestamp on highest-level directories only, not on intermediate directories created as it extracts child files recursively.
- M *maxfiles* (default is unlimited) specifies the maximum number of member files allowed when you use -c to create an HTAR archive.
- o (used only with -x) (default for all nonroot users) causes the extracted files to take on the user and group ID (UID, GID) of the person running HTAR, not those of the original archive. This makes a difference for root users but not for ordinary HTAR users.

- p preserves all UNIX permission fields (on extracted files) in their original modes, ignoring the present UMASK (the default changes the permissions to the local UMASK where HTAR extracts the files). Root users can also preserve the setuid, setgid, and sticky bit permissions with this option.
- S *bufsize* (default is 8 Mbyte) specifies the buffer size to use when HTAR reads from or writes to an HPSS archive file. Here *bufsize* can be a plain integer (interpreted as bytes), an integer suffixed by k, K, kb, or KB for kilobytes, or an integer suffixed by m, M, mb, or MB for megabytes (e.g., 16mb). HTAR's default *bufsize* is picked to maximize its file-transfer performance on the ASCI White machine and -S is intended mostly for LC staff, not ordinary HTAR users (note that HTAR does *not* use the default PFTP block size of 1 Mbyte for storage transfers, nor the 256-kbyte block that is the HPSS default at some nonLLNL HPSS sites).
- T *maxthreads* specifies the maximum number of threads that HTAR will use to copy member files to or from the archive file (default is 15, and HTAR reports the actual number used on each run if you invoke -v or -V). HTAR uses the ratio of buffer size (see -S) to average member file size to estimate how many threads to deploy.
- V (uppercase vee) requests "slightly verbose" reporting of file-transfer progress (often very brief, overwritten messages to the terminal). Do not use with -v.
- v (lowercase vee) requests "very verbose" reporting of file-transfer progress. For each member file transferred to an archive, HTAR prints A (added) and its name on one line; for each member file extracted from an archive, HTAR prints X, its name, and its size on a line, along with a summary of the whole transfer at the end. Do not use with -V.
- w (works only with -x, not with -c) lists (one by one) each member file to be extracted from the archive and prompts you for your choice of confirmatory action, where possible responses are:
- | | |
|--------|---|
| y[es] | extracts the named file. |
| n[o] | skips the named file. |
| a[ll] | extracts the named file and all remaining (not yet processed) selected files too. |
| q[uit] | skips the named file and stops prompting. HTAR ends. |
- Y auto | [*archiveCOS*][:*indexCOS*]
- specifies the HPSS class of service (COS) for each stored archive and its corresponding index file. The default is AUTO, which causes HTAR to use a site-specific COS chosen for archive suitability (at LC, the default COS for HTAR files is 140, which automatically stores a single copy of each archive, regardless of its size). You can

specify a nondefault COS for the archive, the index, or both (e.g., -Y 120:110), but this is usually undesirable except when testing new HPSS features or devices (if your archive size grows to exceed that allowed by a nondefault COS, HPSS will stop the transfer and HTAR will end with an error). Use the special COS 150 (that is, -Y 150) to request dual-copy storage of any mission critical archive of any size for extra safety. Using -Y overrides the HTAR_COS environment variable.

HTAR Examples

Creating an HTAR Archive File

GOAL: To create an HTAR archive file in a subdirectory of your storage home directory and use a filter to install several files within that stored archive.

STRATEGY: (1) One HTAR execute line can perform all of the desired tasks quickly and in parallel:

- The `-cvf` options create (c) an archive, verbosely (v) report the incoming files, and (f) name the envelope file.
- The relative pathname `case3/myproject.tar` locates the archive (`myproject.tar`) in pre-existing subdirectory `case3` of your storage home directory (omitting `case3/` leaves the archive at the top level of your storage home directory). HTAR will not create `case3` itself, however; you must have previously used FTP's `mkdir` option.
- File filter `tim*` selects all and only the files whose names begin with `TIM` (in the directory where you run HTAR) to be stored in the archive.

(2) HTAR opens a preauthenticated connection to your storage (HPSS) home directory and reports its housekeeping activities (very quickly, in lines that overwrite, so you may not notice all of these status reports on your screen).

(3) HTAR creates your requested archive and uses parallel FTP to move your requested files directly into it.

(4) The last incoming file that HTAR reports is always the 256-byte consistency file by which HTAR coordinates your archive with its external index file.

(5) HTAR summarizes the work done (time, rate, amount, thread count), then copies into storage the index file that it made, destroys the local version, and ends.

```
htar -cvf case3/myproject.tar tim*                --- (1)

HTAR: Opening HPSS server connection              --- (2)
HTAR: Getting HPSS site info
HTAR: Writing temp index file to /usr/tmp/aaamva09A

HTAR: creating HPSS Archive file case3/myproject.tar --- (3)
HTAR: a    tim1.txt
HTAR: a    tim2.txt
HTAR: a    tim2a.txt
HTAR: a    tim3.a
HTAR: a    time.txt
HTAR: a    time2.gif
HTAR: a    /tmp/HTAR_CF_CHK_13805_997722535        --- (4)
```

HTAR: Create complete. 29,696 bytes written, ---(5)
max threads: 8
Transfer time: 0.352 seconds (0.084 MB/s)
HTAR: Copying Index File to HPSS...Creating file
HTAR: HTAR SUCCESSFUL

Retrieving Files from within an Archive

GOAL: To retrieve several files from within an existing stored HTAR archive file (without retrieving the whole archive first).

STRATEGY: HTAR does not process metacharacters (file filters such as *) itself, but leaves them for the shell to expand and compare with file names in your *local* directory. Hence, you CANNOT use * to select a subset of already archived files to retrieve. For example, "natural" execute lines

```
htar -xvf case3/myproject.tar time*      [WRONG]
htar -xvf case3/myproject.tar 'time*'    [WRONG]
```

both FAIL to select (and hence to retrieve) any stored files from the MYPROJECTS.TAR stored archive (each yields its own set of error messages). These lines work only accidentally, if you happen to have files with the same name in both your local directory and your stored archive (unlikely except when you are just testing HTAR).

WORKAROUNDS:

(1A) Type the name of each file that you want to retrieve (at the end of the HTAR execute line).

(1B) If you have a long list of files to retrieve, or if you plan to reuse the same retrieval list often, put the list of sought files into a file and use HTAR's -L option to invoke that list. You can use HTAR's -t (reporting) option to help generate that retrieval list by reporting all the files you have archived and then editing that report to include only the relevant file names to retrieve. For instance,

```
htar -tf case3/myproject.tar > hout
grep 'time' hout | cut -c 50-80 > tlist
```

captures the list of all your stored files in the local file HOUT, and then selects just the file names that contain the string TIME for use with HTAR's -L option (here, in local file TLIST).

(1) Once you have laid the groundwork above, a single HTAR execute line can retrieve your specified files quickly and in parallel from within your stored archive:

- The -xvf options request retrieval/extraction (x), verbosely (v) report the retrieved files, and (f) name the target archive.
- The relative pathname case3/myproject.tar locates the archive (myproject.tar) in pre-existing subdirectory case3 of your storage home directory.

- The explicit file list (1A) or name-containing file (1B) selects all and only the files that you want (here, those whose names begin with TIME, a subset of all files stored in this archive in the previous example).

(2) HTAR opens a preauthenticated connection to your storage (HPSS) home directory and reports its housekeeping activities (very quickly, in lines that overwrite, so you may not notice all of these status reports on your screen).

(3) HTAR uses its external index to locate in the archive the (two) specific files that you requested and then it transfers them by parallel FTP to your local machine *without* retrieving the whole archive file.

(4) HTAR summarizes the work done (time, rate, amount) and then ends.

```
htar -xvf case3/myproject.tar time.txt time2.gif          --- (1A)
OR
htar -xvf case3/myproject.tar -L tlist                    --- (1B)

HTAR: Opening HPSS server connection                      --- (2)
HTAR: Reading index file
HTAR: Opening archive file

HTAR: Reading archive file                                --- (3)
HTAR: x  time.txt, 1085 bytes, 4 media blocks
HTAR: x  time2.gif, 3452 bytes, 8 media blocks

HTAR: Extract complete. total bytes read:                --- (4)
      6,144 in 0.088 seconds (0.070 MB/s)
HTAR: HTAR SUCCESSFUL
```

Rebuilding a Missing Index

GOAL: To rebuild the missing index file for a stored HTAR archive file and thereby (re)enable blocked access to the files within it (and extract some).

STRATEGY:

- (1) You try to retrieve all files (-xvf) from the HTAR archive myproject.tar in the case3 subdirectory of your storage home directory.
- (2) But HTAR cannot find the external index file (here, called myproject.tar.idx) for this archive, and it returns a somewhat cryptic error message, retrieves no requested files, and ends. (File myproject.tar.idx may have been moved, renamed, or accidentally deleted from storage.)
- (3) So you execute HTAR again with the special action -X (uppercase, not lowercase, eks) to request rebuilding the external index for the (same) disabled archive.
- (4) HTAR opens a preauthenticated connection to your storage (HPSS) home directory, locates the archive in subdirectory case3, scans (but does not retrieve) its contents, and thereby creates a new myproject.tar.idx file (temporarily on local disk, then moved to the same storage directory as the archive file that it supports). HTAR ends.
- (5) Now you again try your original (1) file-retrieval request.
- (6) HTAR opens a preauthenticated connection to your storage (HPSS) home directory and reports its housekeeping activities (very quickly, in lines that overwrite, so you may not notice all of these status reports on your screen).
- (7) HTAR uses its (newly rebuilt) external index to locate the files within the archive and transfers them by parallel FTP to your local machine (it transfers all of them because there is no *filelist* on the execute line).
- (8) HTAR summarizes the work done (time, rate, amount) and then ends.

```
htar -xvf case3/myproject.tar --- (1)
```

```
HTAR: Opening HPSS server connection
HTAR: Getting HPSS site info
ERROR: Received unexpected reply from server: 550 --- (2)
ERROR: Error -1 getting Index File attributes...
HTAR: HTAR FAILED
###WARNING htar returned non-zero exit status.
          72 = /usr/local/bin/htar.exe...
```

```
htar -Xf case3/myproject.tar --- (3)
```

```
HTAR: Opening HPSS server connection --- (4)
HTAR: Reading archive
HTAR: Copying Index File to HPSS... creating file
HTAR: HTAR SUCCESSFUL
```

```
htar -xvf case3/myproject.tar --- (5)
```

```
HTAR: Opening HPSS server connection --- (6)
HTAR: Reading index file
```

```
HTAR: Opening archive file

HTAR: Reading archive file ---(7)
HTAR: x tim1.txt, 3503 bytes, 8 media blocks
HTAR: x tim2.txt, 4310 bytes, 10 media blocks
HTAR: x tim2a.txt, 5221 bytes, 12 media blocks
HTAR: x tim3a., 5851 bytes, 13 media blocks
HTAR: x time.txt, 1085 bytes, 4 media blocks
HTAR: x time2.gif, 3452 bytes, 8 media blocks

HTAR: Extract complete. total bytes read: ---(8)
      28,160 in 0.141 seconds (0.200 MB/s)
HTAR: HTAR SUCCESSFUL
```

Specifying Very Many Files

GOAL: To specify a very large number of input files yet avoid an HTAR command line too long for the shell to accept.

STRATEGY: The input-line-too-long problem is analyzed and explained in the HTAR [Limitations](#) (page 15) section above. The best solution is to keep in a separate directory all and only the intended input files, and then specify that directory's name on HTAR's execute line (for recursive processing). But if you failed to take that precaution, you can work around the problem of having too many file names for the shell to accept by using the UNIX FIND utility as shown here.

(1) Run FIND to select the files that you want (here, those whose names begin with T) in a way that processes the file list internally, not by the shell. Collect the file names in an intermediate file (here, called TFILES). The [Limitations](#) (page 15) section above explains why FIND's -EXEC option will *not* do the job here.

(2) Run HTAR to build the stored archive that you want (-cf) using the -L option to read in the list of input files (TFILES) that FIND compiled.

(3) Without -v, HTAR shows no verification but does copy its index file to storage when the archive is done.

(4) Remove the (long) list of names, no longer needed.

```
find . -name 't*' -print > tfiles --- (1)
```

```
htar -cf test.tar -L tfiles --- (2)
```

```
HTAR: Opening HPSS server connection
HTAR: creating HPSS archive file test.tar
HTAR: Copying index file to HPSS...creating file --- (3)
HTAR: HTAR SUCCESSFUL
```

```
rm tfiles --- (4)
```

Archiving Between Nonstorage Machines

GOAL: To put files from a directory local to one LC production machine into an HTAR archive file that resides in a directory local to another LC production machine (*not* in HPSS storage as usual).

STRATEGY: Your common home directory is already cross-mounted on every LC production machine (and /nfs/tmp is cross-mounted on many machines), so you never need to use HTAR to transfer home files "between machines." And if you want an HTAR archive deposited on the *same* machine where HTAR runs (instead of in storage), use HTAR's -E option. You only need to follow the steps below if you really need to transfer files from one machine's *local* directory (such as /usr/tmp) to or from an archive file in another machine's *local* directory. Use -F to specify the nonstorage archive host.

(1) To create (-c) an HTAR archive on another machine (instead of storage), run HTAR on the machine and in the directory where the files to be archived *reside* (because HTAR makes an archive using PUTs). For example:

SOURCE:	SINK:
ILX1	GPS17
/var/tmp	/usr/tmp
file1...n	myarchive.tar
-----create--->>>	myarchive.tar.idx
[HTAR runs here]	

Use -F to specify the sink machine (where the archive and its index file will go, here GPS17), and use -f to specify the full pathname of the archive on that sink machine. Note that HTAR still contacts the HPSS server to log this transaction, even though the archive is not stored in HPSS.

(2) To extract (-x) file1 from an HTAR archive on another machine (instead of from storage), run HTAR on the machine and in the directory where the files to be extracted should *arrive* (because HTAR uses GETs to extract files). For example:

SINK:	SOURCE:
ILX1	GPS17
/var/tmp	/usr/tmp
file1	myarchive.tar
<<<---extract----	myarchive.tar.idx
[HTAR runs here]	

Use -F to specify the source machine (where the archive and its index file reside, here GPS17), and use -f to specify the full pathname of the archive on that source machine. Note that HTAR still contacts the HPSS server to log this transaction, even though the archive is not stored in HPSS.

```
htar -c -f /usr/tmp/myarchive.tar -F gps17 *          ---(1)
```

```
HTAR: Opening FTP server connection
HTAR: Opening HPSS server connection
```



```
    Creating FTP archive file /usr/tmp/myarchive.tar
    Pass 2: adjusting local index file entries
    Copying index file to remote host...creating file
HTAR: HTAR SUCCESSFUL

htar -x -f /usr/tmp/myarchive.tar -F gps17 file1      ---(2)

HTAR: Opening FTP server connection
HTAR: Opening HPSS server connection
      Reading index file
      Opening archive file
      Reading archive file
HTAR: HTAR SUCCESSFUL
```

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.

(C) Copyright 2003 The Regents of the University of California. All rights reserved.

Keyword Index

To see an alphabetical list of keywords for this document, consult the next section (page 36).

Keyword	Description
<u>entire</u>	This entire document.
<u>title</u>	The name of this document.
<u>scope</u>	Topics covered in this document.
<u>availability</u>	Where HTAR runs.
<u>who</u>	Who to contact for assistance.
<u>introduction</u>	General HTAR overview, analysis.
<u>htar-role</u>	Goals, scope, performance of HTAR.
<u>htar-files</u>	Three key HTAR files diagrammed.
<u>tar-comparison</u>	TAR and HTAR features compared.
<u>htar-usage</u>	How to use HTAR.
<u>execute-line</u>	Required syntax, features, defaults.
<u>htar-errors</u>	Common errors conditions, warnings.
<u>limitations</u>	Known HTAR limitations, work-arounds.
<u>environment-variables</u>	Env. variables used by HTAR.
<u>options</u>	HTAR options grouped, explained.
<u>action</u>	HTAR's action options (1 reqd).
<u>archive</u>	HTAR's archive option (always reqd).
<u>control</u>	HTAR's control options.
<u>examples</u>	Annotated sample HTAR sessions.
<u>create-archive</u>	How to make an HTAR archive.
<u>retrieve-files</u>	How to extract HTAR files.
<u>rebuild-index</u>	How to rescue a lost HTAR index.
<u>many-files</u>	Specifying very many input files.
<u>find-input</u>	Specifying very many input files.
<u>between-machines</u>	Archiving between NONstorage machines.
<u>index</u>	The structural index of keywords.
<u>a</u>	The alphabetical index of keywords.
<u>date</u>	The latest changes to this document.
<u>revisions</u>	The complete revision history.

Alphabetical List of Keywords

Keyword	Description
-----	-----
<u>a</u>	The alphabetical index of keywords.
<u>action</u>	HTAR's action options (1 reqd).
<u>archive</u>	HTAR's archive option (always reqd).
<u>availability</u>	Where HTAR runs.
<u>between-machines</u>	Archiving between NONstorage machines.
<u>control</u>	HTAR's control options.
<u>create-archive</u>	How to make an HTAR archive.
<u>date</u>	The latest changes to this document.
<u>entire</u>	This entire document.
<u>environment-variables</u>	Env. variables used by HTAR.
<u>examples</u>	Annotated sample HTAR sessions.
<u>execute-line</u>	Required syntax, features, defaults.
<u>find-input</u>	Specifying very many input files.
<u>htar-errors</u>	Common errors conditions, warnings.
<u>htar-files</u>	Three key HTAR files diagrammed.
<u>htar-role</u>	Goals, scope, performance of HTAR.
<u>htar-usage</u>	How to use HTAR.
<u>index</u>	The structural index of keywords.
<u>introduction</u>	General HTAR overview, analysis.
<u>limitations</u>	Known HTAR limitations, work-arounds.
<u>many-files</u>	Specifying very many input files.
<u>options</u>	HTAR options grouped, explained.
<u>rebuild-index</u>	How to rescue a lost HTAR index.
<u>retrieve-files</u>	How to extract HTAR files.
<u>revisions</u>	The complete revision history.
<u>scope</u>	Topics covered in this document.
<u>tar-comparison</u>	TAR and HTAR features compared.
<u>title</u>	The name of this document.
<u>who</u>	Who to contact for assistance.

Date and Revisions

Revision Date -----	Keyword Affected -----	Description of Change -----
17Nov03	<u>htar-role</u> <u>limitations</u> <u>execute-line</u> <u>action</u>	Size and speed details updated. Size and speed details updated. General transfers with -F clarified. General transfers with -F noted.
23Jul03	<u>between-machines</u> <u>index</u> <u>examples</u> <u>tar-comparison</u> <u>limitations</u> <u>control</u> <u>introduction</u>	New section on nonstorage archiving. New keyword for section. More verbose, explicit HTAR output. Details updated for new features. Only passwordless FTP servers allowed. -F, -H, -M control options added. Class of service clarified.
13May03	<u>availability</u> <u>htar-role</u> <u>execute-line</u>	HTAR now runs under Linux/CHAOS. Use NETMON for performance data. Consistency file location under Linux.
24Jun02	<u>htar-role</u> <u>execute-line</u> <u>limitations</u> <u>control</u>	Size and copy issues clarified. Y option added. Single-copy default now, dual with Y. Maximum file size clarified. -Y 150 for dual copy noted.
01May02	<u>htar-errors</u> <u>limitations</u> <u>control</u> <u>retrieve-files</u>	Retrieval with filters problem noted. Retrieval with filters problem noted. -L use expanded. Example now treats filter problem.
17Oct01	<u>options</u> <u>index</u>	-S and -h roles clarified. Options section subdivided. New keywords for subsections.
20Aug01	entire	First edition of HTAR manual.

TRG (17Nov03)

UCRL-WEB-200720

Privacy and Legal Notice (URL: <http://www.llnl.gov/disclaimer.html>)

TRG (17Nov03) Contact on the OCF: lc-hotline@llnl.gov, on the SCF: lc-hotline@pop.llnl.gov